

WHAT MAKES VICARIOUS FUNCTIONING WORK? EXPLORING THE GEOMETRY OF HUMAN-TECHNOLOGY INTERACTION

Asaf Degani
NASA Ames Research Center
Moffett Field, CA

Michael Shafto
NASA Ames Research Center
Moffett Field, CA

Alex Kirlik
University of Illinois
Champaign, IL

*The essence of modeling lies in establishing relations between pairs of
system descriptions.* (Zeigler, Praehofer, & Kim 2000; p. 295)

INTRODUCTION

In this chapter we provide a Brunswikian perspective on human interaction with everyday technologies such as traffic lights, automotive devices (e.g., warning systems), and also advanced technologies such as flight control systems in modern airliners. We apply this perspective toward suggesting a framework for evaluating interface designs and for ultimately improving the usability, robustness, and effectiveness of a range of interactive technologies.

Today's automated systems, such as modern commercial "glass cockpit" aircraft, afford the user various levels, or "modes," of interaction, ranging from fully manual to fully automatic. These modes provide the pilot with various control strategies to achieve a given goal. In all automated control systems, including those found in cars, ships, and aircraft, the control modes are discrete whereas the behavior of the controlled system (e.g., the aircraft) is continuous. In commercial aviation, the pilots' task of coping with the mapping between discrete mode changes and dynamic, continuous changes such as altitude, heading, and speed is challenging. Incident and accident data show a strong relationship between environmental demands (e.g., air traffic control clearances), mode-selection strategies, interface design, and operational problems (Degani, 1996, see Chapter 7).

We believe that a deeper understanding of the nature and implications of the relationships between the demands of the operational environment (e.g., air traffic control), the

physical space in which performance occurs (e.g., the airspace), the technology that's employed (the automation and its interfaces), and, finally, human cognition are important for enhancing human interaction with the semi-automated systems of today and, hopefully, the autonomous systems (e.g., planetary rovers) of the future.

VICARIOUS FUNCTIONING, A.K.A. "PURPOSIVE BEHAVIOR"

"There is a variety of 'means' to each end, and this variety is changing, both variety and change being forms of vicarious functioning" (Brunswik, 1952, p. 18). In today's cognitive parlance, vicarious functioning might be glossed as *flexible, goal-oriented behavior*. Wolf (1999) elaborates on the central role of this type of adaptive behavior in Brunswik's work by noting that Brunswik emphasized the adaptive relations between an organism and its environment. His analysis of the environment revealed it to be stochastic, dynamic, and non-repeating; at once ambiguous and partially redundant. In response, adaptive behavior requires what Wolf terms a "virtuosity of replacement," or an ability to select from a wide repertoire of adaptations in response to the dynamic structure of the environment at each point in time.

As a synonym for *vicarious functioning*, Brunswik himself (1952, p. 16) used the more transparent phrase *purposive behavior*. He noted that without a goal or purpose, behavior itself is hard to define: He agreed with E. G. Boring that nothing would make a robot seem more human than "an ability to choose one means after another until the goal is reached" (Brunswik, 1952, p. 17). Wolf (1999) notes the similarity of Brunswik's concept of vicarious functioning to a currently popular approach in AI, i.e., "reactive" (incremental, least-commitment) planning:

"According to Brunswik it is typical for humans to make use of alternatives, to commit only provisionally in order to keep possibilities for revision open [cf. Connolly, 1999]. For human perception—or more generally cognition—as well as for human action, it is necessary to cope with inconsistent, unexpected, incomplete, and imperfect events."

This emphasis on incremental opportunism and robustness has connections with other distinctive characteristics of Brunswik's psychology: the entire program of probabilistic functionalism, the lens model, and the requirement to conduct research in representative environments (Brunswik, 1956). If dynamic and stochastic ecological processes are abstracted away, leaving only the bite-board and the response-key, then neither perception nor the control of action can exhibit their evolved relations. As Allen Newell memorably put it, "If you study simple systems, you will learn a lot about simple systems."

A FIELD STUDY OF HUMAN-AUTOMATION INTERACTION IN COMMERCIAL AVIATION

Degani (1996) made cockpit observations of pilots' interactions with the automatic flight control system of the Boeing 757/767 aircraft during 60 flights (cf. Casner 1994; this volume). During the flights, every observable change in the aircraft's control modes, either manually initiated, (e.g., the pilot selected a new mode) or automatically initiated, (e.g. an automatic mode transition) was recorded, along with all the parameters relating to

the flight control system status (e.g., waypoints and altitude values selected by the pilot). Likewise, every observable change in the operating environment (e.g., a new ATC instruction, switching from one ATC facility to another) was recorded, along with other variables such as the aircraft's altitude, speed, and distance from the airport. In a way, it was like taking a snapshot of every change that took place in the cockpit. Overall, the dataset consisted of 1,665 such snapshots. Each snapshot consisted of 18 categories describing the status of the automatic flight control system, and 15 categories describing the operational environment. Data analysis and presentation is discussed below.

Analysis, Visualization, and Interpretation

Jha and Bisantz (2001; also see Jha & Bisantz this volume) noted that the multivariate lens model can be extended to the analysis of categorical judgments. Methods for the analysis of multivariate categorical data were integrated and generalized during the 1980s and 1990s, by relating multivariate analysis to graphical algorithms (De Leeuw & Michailidis, 2000). This line of work unifies principal components analysis, canonical correlation analysis, and other types of multivariate analysis, and extends their coverage to include categorical data.

To analyze the field data described above, we used one version of this type of approach, canonical correlation analysis, to quantify the relations between patterns of environmental variables and patterns of mode selections (see Degani, 1996, and Shafto, Degani, & Kirlik, 1997, for details). We considered the environmental patterns to be the *independent* variables (X) and the automation mode selection patterns to be the *dependent* variables (Y). Conceptually speaking, this analysis correlates two multivariate patterns in the same way bivariate correlation measures the relationship between two single (X, Y) variables.

Due to the high dimensionality of our dataset, we explored several graphical methods to help us understand and communicate the relationships between the two multidimensional patterns found by the canonical correlation. One of the most helpful suggestions we found was due to Cliff (1987), advocating presentation of structure correlations rather than weights. Structure correlations are the correlations of the X canonical variate with each of the original independent variables, and of the Y canonical variate with each of the original dependent variables. In this way, the canonical variates can be interpreted in terms of observed variables.

To display the independent X (environmental) and dependent Y (technology use) patterns, we developed a visual display we termed a "heliograph." In this sunburst-like display (see Figure 1), the relative sizes of the structured correlations are indicated by the lengths of the bars extending outward or inward, both indicating relations between features of the ecology and features of technology use.

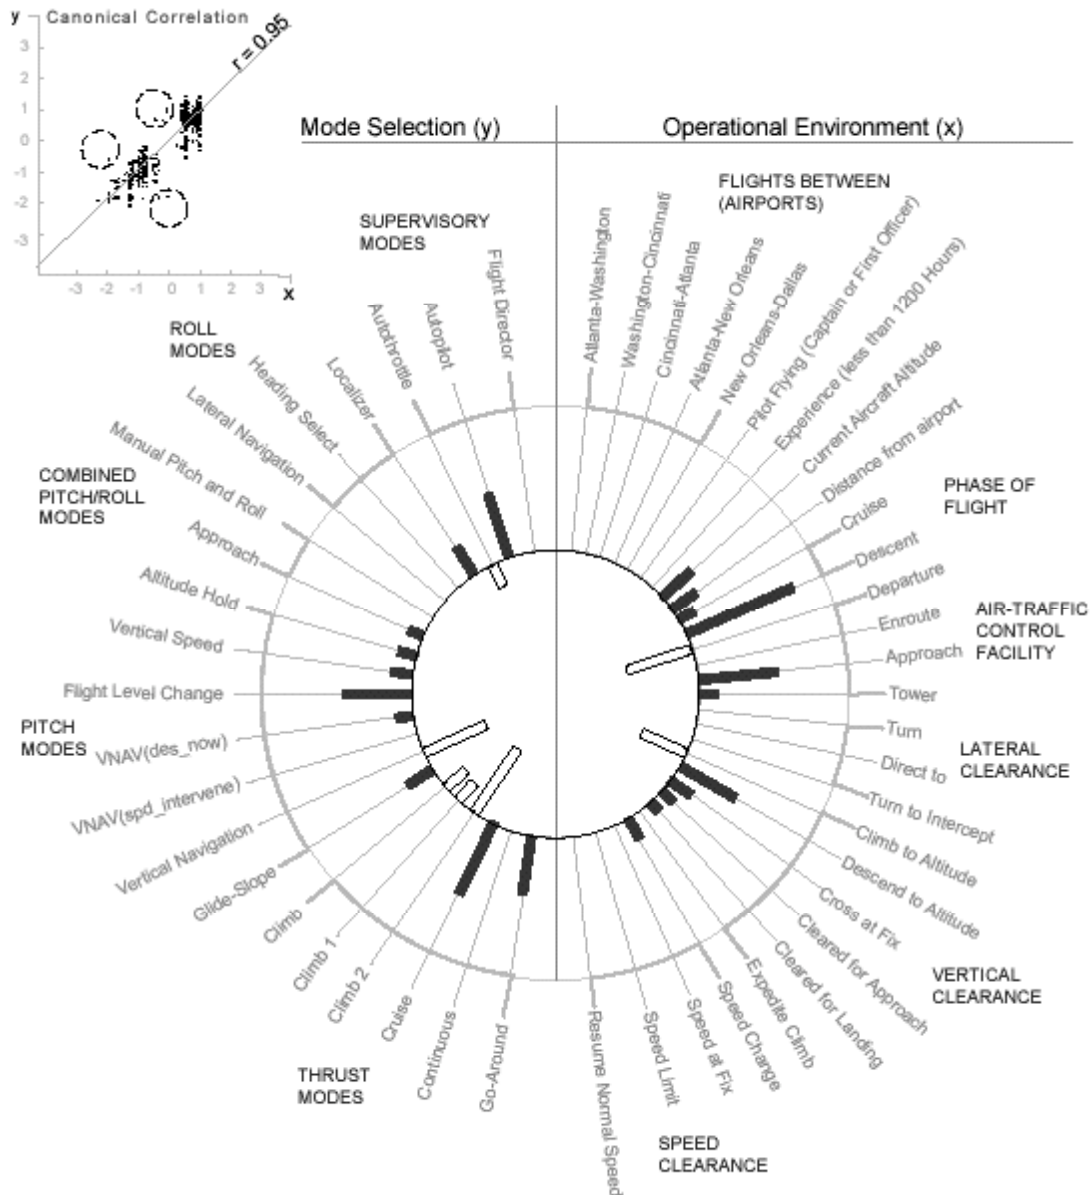


Figure 1. Heliograph Showing the Patterns ($r = .95$) Between the Features of the Operational Environment (x) and Pilots' Mode Selections (y)

Figure 1 shows two strong patterns between the demands of the operational environment (X), and mode selections (Y): The first pattern (dark bars), shows that when aircraft altitude is above average (of 13,000 feet), during the “descent” (Phase of flight), and while under “approach” (ATC Facility), and in response to a “descend to altitude” clearance—flightcrews usually engage “Flight Level Change” (Pitch) mode, cruise (Thrust) mode, and engage the autopilot. The behaviors can all be classified as engaging “supervisory” modes of automation, in that automation has direct control over the aircraft, while the flight crew mainly supervises, monitors, and intervenes only when necessary (see Sheridan, 1997, on “supervisory control”).

The second pattern (white bars) shows that while under “departure” (ATC Facility) control and in response to a “climb to altitude” clearance, flightcrews usually engage “Vertical Navigation” (Pitch) mode, climb (Thrust) mode, and the autothrottles. In addition to the engaged modes, it is also important to note in Figure 1 what is not engaged. In the first pattern (dark bars), note that the “autothrottles” did not appear, meaning that many pilots do not engage the autothrottles during the descend/approach phase (rather, they manually move the throttles). As for the second pattern (white bars), note that the autopilot is not always used during departure and initial climb-out and that many pilots prefer to hand fly the aircraft during this phase of flight. Finally, the plot in the upper left corner of the graph is a conventional scatter diagram showing the relationship between the two composite variables (X and Y), plotted here in standard units, suggesting a strong ($r=0.95$) overall relationship between the patterns.

It is possible to use canonical correlation to identify additional patterns in data sets such as these, and the reader can refer to Degani (1996) for three other heliographs based on the same data set demonstrating additional environmental-automation usage patterns with correlations of $r=0.88$, $r=0.81$, and $r=0.72$. Overall, the canonical correlation analysis produced eight meaningful, consistent, and statistically independent patterns that were later corroborated by expert pilots. Two of the empirically derived patterns actually turned out to mirror standard operating procedures at this particular airline.

In addition to identifying patterns, canonical correlation can be used to help the analyst define, in a data-driven way, the most important environmental cues that affect judgment and action selection on the user’s side (in this case, mode selection). Equally as important are deviations from any central tendencies in these patterns. Such deviations (outliers) from a standard pattern of human-automation interaction derived from data across many users can mean two things: either an unusual (yet safe) mode selection, or a dangerous (and potentially unsafe) mode selection.

For example, one flight in our data set was an obvious outlier, falling 7 standard deviations from the mean on one of the spatial dimensions. Inspection of the data revealed that the autothrottles were not used throughout the flight. The reason for this was that the aircraft had been dispatched with inoperative autothrottles, and hence they were not used throughout the flight. This situation was definitely unusual, but not unsafe, because the crew knew of the inoperative autothrottles well before the flight, and planned their flight accordingly.

As for unsafe outliers, two flights were observed falling 3 standard deviations from the mean. These flight crews used the “Flight Level Change” mode while making a final approach for landing during what is termed a “back-course localizer” approach to a runway. This situation, in which the glideslope instrument system is unavailable, is quite a rare occurrence (almost all runways that are used by large aircraft do indeed have glideslope instrumentation, yet sometimes the system malfunctions or is shutdown for maintenance). The selection of the “Flight Level Change” mode in this particular situation is potentially dangerous during final approach and many airlines have procedures in place to warn pilot against using it at low altitude. Indeed, its use at very low altitudes has contributed to several near fatal incidents and two accidents (Indian Court of Inquiry, 1992; Ministry of Planning, Housing, Transport and Maritime Affairs,

1989). Usually, flight crews use the “Flight Level Change” mode during descent, and then switch to “Glide Slope” mode on final approach. But on these two occasions during the field study, because there was no glideslope instrument (and therefore no cue to prompt crews to switch modes), flight crews kept on descending using the “Flight Level Change” mode, which, as mentioned earlier, is not safe for making a final approach.

We believe that the above examples of unsafe mode usage represent instances of a broader class of human-technology interaction problems associated with discretely mediated interaction with, and control over, variables that are highly dynamic. Presenting a system operator, such as a pilot, with an abstracted, discrete suite of action opportunities, and an abstracted, solely discrete display of automation operation (e.g., mode settings) can lead to situations where the controlled system drifts away into unsafe regions where the fixed mode setting is no longer appropriate. This discrete approach to automation and interface design, we believe, short-circuits some of the naturally evolved psychological mechanisms supporting vicarious functioning in an ecology that is, or at least once was, inherently continuous. With the result, we suggest, of undermining some of the basic mechanisms of adaptive behavior itself. We now turn to examining and elaborating this diagnosis and hypothesis in more detail.

WHAT KINDS OF ENVIRONMENTS SUPPORT VICARIOUS FUNCTIONING?

Understanding, and especially predicting, human adaptation to dynamic environments is by no means simple nor straightforward (Bullock & Todd, 1999; Kirlik, in press). How, then, should we think about the design of digital and automated control systems so as to meet the computational requirements posed by various tasks in these environments (e.g., aircraft navigation) on the one hand, and also leverage, if not short-circuit, our evolutionarily acquired resources for flexible, goal-directed behavior on the other? One approach that we believe is promising is to consider the *physical space* in which behavior occurs, the *technological spaces* that increasingly characterize the ecology of modern life, and the *psychological space* in which adaptive mechanisms operate.

Mapping out these spaces, and most importantly the relations among them, is what we allude to in our title as characterizing and analyzing “the geometry of human-automation interaction.”

Physical Space

Natural, physical space provides the strongest set of constraints on design, since we are well equipped to adapt to constraints imposed by, inherently continuous, physical space. Rarely do we try to walk through walls. Any discussion of physical space at the scale of the human ecology starts with the four-dimensional trajectories to which we are comfortably adapted. As long as four-dimensional Euclidean space is the only option, we can rely on the typically fluent and highly adaptive perception and action mechanisms that are keyed to getting around in this space; these are evolution’s legacy.

Technological Space

In contrast, when we consider the computer-based artifacts increasingly present in our ecology, design options mainly span or sample discrete spaces (e.g., see Pirolli, this

volume). The reason for this lies in the logic governing the behavior of any computer system rooted in the discrete, finite-state machine (FSM) formalism, the von Neumann architecture, and the Turing-inspired, algorithmic approach to software specification. Additionally, in some automated control systems, we encounter hybrid spaces which harbor both discrete logic (modes) and continuous parameters (such as speed and flight path angle) that are based on laws of physics. We also encounter hybrid spaces in computer networks, in which continuous variables like task priority and processing time must also be considered to achieve robust communication and coordination (Lowe, 1992; Roscoe, 1998; Schneider, 2000). The geometry of the computational ecology is complex indeed. As we will suggest in the following, this complexity becomes mirrored in the cognitive activities necessary to adapt to, and navigate through, these discrete and hybrid ecologies.

Psychological Space

We now turn our attention to more distinctly psychological issues, relying on slightly more metaphorical notions of space and geometry. A large and growing literature supports the conclusion that psychological “space” also has both continuous and discrete dimensions. Continuous psychological processes and representations mirror the structure of physical space, in that they have few dimensions, are strongly continuous, and are strongly metric. In contrast, symbolic processes based in language and logic, which serve as the basis for analytical cognition, are largely discrete, are weakly metric or non-metric, and are typically discontinuous.

Kahneman (2003; see also Hammond, 1996, Hastie & Dawes, 2001, p. 4; and Sun, 2000) contrasts between his notions of “System 1” (our “continuous”) and his “System 2” (our “discrete”) psychological activities:

“The operations of System 1 are typically fast, automatic, effortless, associative, implicit (not available to introspection), and are often emotionally charged; they are also governed by habit and are therefore difficult to control or modify. The operations of System 2 are slower, serial, effortful, more likely to be consciously monitored and deliberately controlled; they are also relatively flexible and potentially tend to disrupt each other, whereas effortless processes neither cause nor suffer much interference when combined with other tasks.” (p. 699)

In their analysis of dual-process theory (e.g., Kahneman’s System 1 and System 2), Norenzayan, Smith, Kim, & Nisbett (2002) suggest that intuitive cognition (our “continuous,” Kahneman’s “System 1”) tends to dominate over formal cognition (our “discrete,” Kahneman’s “System 2”), although both systems are typically active simultaneously:

“In recent years, a growing number of research programs in psychology have examined these two cognitive systems under the rubric of ‘dual process’ theories of thinking....These two cognitive systems coexist in individuals, interact with each other in interesting ways, and occasionally may be in conflict and produce contradictory inferences.... intuitive reasoning [continuity] tends to dominate, but the relative dominance can be modulated by a variety of factors.” (pp. 654-655)

Current wisdom suggests that people can operate in either of these two contrasting modes of cognition. A pioneer in theory along these lines, Hammond (1996), motivated by Brunswik's (1956) initial observations on the distinctions between perceiving and thinking, has developed an even more sophisticated theory in which people can operate, and even oscillate between, various points along a "cognitive continuum" with both analytical and intuitive poles. Researchers such as Gigerenzer and Goldstein (1996), Kirlik (1995), Klein (1999), Norenzayan et al. (2002), Reason (1990), and Rouse (1983), among others, have all suggested that people tend to rely on intuitive as opposed to analytical cognition when possible, although this preference can be mitigated by a variety of factors (e.g., the demand for accountability or to justify one's actions, necessitating accessibility to the processes of thought, and thus more analytical cognition).

In short, we conceive of skilled, fluent, and robust adaptive behavior as relying heavily on relatively resource-unlimited, intuitive cognition, which is most naturally supported (as we will suggest in more detail in the following) by continuous ecologies. More resource-intensive analytical cognition, in contrast, functions mainly to monitor the ongoing success of intuitive behavior, and to intervene when necessary or demanded, given that time and cognitive resources are available to do so (Kahneman, 2003). We next discuss the crucial role played by continuity in supporting the intuitive, adaptive mode behavior Brunswik described as "vicarious functioning."

ECOLOGICAL CONTINUITY: A KEY RESOURCE FOR VICARIOUS FUNCTIONING

The most obvious source of the dominance or preference for cognition that is based on continuous spaces and intuitive, or System 1, activities is their evolutionary priority. These activities are responsible for coordinating perception and action in the physical world, and as such, perception is adapted to the regularities of the physical environment (Barlow, 2001; see also Hubbard, 1999; Shepard, 1999). Experiments by Shepard and others have shown, for example, that the perceived trajectories of objects seen only in successive snapshots obey the kinematics of actual objects moving through physical space, as though physical laws had been internalized. Based on a sustained program of research in this tradition, Shepard (1999) concluded that "...objects support optimal generalization and categorization when represented in an evolutionarily shaped space of possible objects as connected regions with associated weights determined by Bayesian revision of maximum-entropy priors [probabilities]."

The "connected regions" to which Shepard alludes echo our own emphasis on the importance of continuity to cognition and adaptive behavior generally. Shepard's mention of "Bayesian revision" of probabilities points to Brunswik's, and our own, observation that many environments often possess irreducible uncertainty. As such, the human is required to adapt to the world in a statistical, rather than a deterministic sense. James. J. Gibson titled his 1957 *Contemporary Psychology* review of Brunswik's probabilistic functionalism as "Survival in a World of Probable Objects" (Gibson 1957/2001; Kirlik, 2001; Cooksey, 2001). This title suggests that Gibson's gloss of Brunswik's approach to describing the demands of with adapting to an uncertain world was very much in the spirit of Shepard's, presented some 40 years later.

In the spirit of Shepard's research, Barlow (2001) has noted, for example, that there is widespread agreement that the perceived "trajectory of an object is a joint function of perceptually sampled data and of the bias that is intrinsic to psychological space." Perception can thus be regarded as a process at the intersection of internalized, environmental regularities (via evolution, development, and experience) and the information available from the external, currently present environment. Barlow pointed out that it was Brunswik who first suggested that:

"The laws governing grouping and segregation of figure from ground were more than empirical facts about perception: they were rules for using statistical facts about images to draw valid inferences from the scene immediately before the eyes. [Brunswik] pointed out that two perceptual samples having similar local characteristics are likely to be derived from the same object in the external world. Therefore, it is adaptive to have a built-in heuristic bias that they are from the same object."

Barlow (2001) also reviewed recent neurophysiological results which further confirm that the natural ecology's reflection: a four-dimensional, strongly continuous, metric space is "wired into" the central nervous system at a fundamental level. This body of research lends support to the lens model's basis in principles of environmental-psychological symmetry.

Additionally, Barlow has noted the existence of "neurophysiological mechanisms that exploit the redundancy of sensory messages resulting from statistical regularities of the environment," lending additional support to Brunswik's view that vicarious functioning, exploiting these ecological redundancies, may very well play the role of the "backbone of stabilized achievement" (Brunswik, 1956, p. 142).

The question remains, however, as to why a preference for low-dimensional, continuous representations should extend beyond basic perception and action, and also into the cognitive realm, "beyond the information given" (Bruner, 1973). The answer, we believe, lies in the advantages of continuous representations for supporting learning and efficient information processing, advantages to which Shepard (1999, quotation above) alluded.

As demonstrated by the groundbreaking research of Landauer et al. (1997), discrete, symbolic representations such as natural language are not necessarily separable from Kahneman's "System 1" or Hammond's "intuition." Rather, as Landauer et al.'s work demonstrates, a ubiquitous learning process exists to convert the discrete to the continuous. This adaptive process automatically converts symbolic perceptual information, such as text, into continuous, spatial cognitive representations, tends to produce and strengthen constraints on inference, and tends to reduce the dimensionality of the discrete input space. By "strengthening constraints," we refer to a process of mapping discrete or categorical representations to continuous metric spaces, and mapping weakly continuous spaces into strongly continuous ones.

To take a more concrete example, Landauer, Laham, Rehder, and Schreiner (1997) have shown that a linear neural network, operating on a corpus of sequenced discrete tokens (words, for example), will induce a multidimensional continuous space summarizing the tokens and their associated contexts. The space is bootstrapped from the underlying

continuity of time, which enforces the basic sequential ordering of the tokens. On the basis of the induced continuous space, a range of abstract, *symbolic* tasks can be performed at levels of achievement equal to or better than expert human performance. Examples include matching documents of similar meaning but with no shared words, while rejecting documents with shared words but having dissimilar meaning; choosing correct answers on standardized vocabulary tests; responding to semantic priming in laboratory experiments; learning from instructional texts; interpreting novel analogies and metaphors, and, in many instances, grading student essays as well as human graders. Along similar lines, McGreevy and Statler (1998), using a different but related algorithmic technique, have demonstrated that the interpretation and comparison of discrete and symbolic verbal accounts of aviation incidents reports could be enhanced by translating these verbal data into continuous spatial representations.

The Unique Benefits of Continuous Spaces

At a high level we have already discussed some of the benefits that continuous, metric spaces confer upon adaptation. The availability of convergent, recursive learning processes (Landauer et al., 1997) and the compression of information made possible by continuous, metric representations can, in many cases, offset the disadvantages of applying continuous operations to problems that could be more ideally addressed by discrete, logical, and symbolic activities. On the down side, Kahneman (2003) provides a thorough discussion of these potential disadvantages (see also Hastie & Dawes, 2001; Freed & Remington, 1998). On the up side, the recursive learning and information compression enabled by continuous representations point to the most obvious (and some might say, the only) information-processing capabilities of neural networks. This property of intrinsic, environmental, or automatically-derived cognitive continuity, tacitly presumed in so many studies and models of learning and adaptation, is so ubiquitous that it is easily overlooked. Yet, continuity confers an enormous range of functional advantages enabling purposive behavior and adaptation in an environment that is dynamic and unpredictable.

Most importantly, continuity supports *approximation* and *convergence*. These properties enable statistical or neural-network based learning and generalization. In contrast, discrete spaces, characteristic of the interfaces of many everyday technologies (e.g., VCRs, cell phones), do not support efficient learning and generalization, nor do they support flexible goal-directed behavior. Why? Simply put, because one discrete state is no more or less like any other discrete state. As such, unless an interface designer provides cues (e.g., proximity, color or shape coding, hierarchical menu structuring), to *explicitly* support inference, there is typically little support for generalizing what is known about any one discrete state (display or control) to any other.

In the natural ecology at the human scale, continuity as a resource for learning and generalization is crucial to guarantee the convergence of adaptive solutions to problems of many types. For example, if we are practicing our piano skills and our teacher tells us that we played a certain note too softly in one case, and too loudly on our next attempt, we can infer that if we play the note at an intermediate level of force on our third try that we are at least *likely* to obtain a better outcome. Contrast this case with trying to adapt to a discrete space, such as the state-space of a digital wristwatch. Given that you know the

result of pushing two of the three buttons on the watch, what can you learn about what is likely to happen if you push the third? Nothing.

In some cases, of course, a continuous, multidimensional space simply cannot accurately represent all the necessary distinctions required to achieve perfect adaptation, as is sometimes the case in understanding natural language. Yet, as Dawes (1979) has suggested in his classic article, “The robust beauty of improper linear models,” the degree of “meaning” that is lost by moving from discrete and symbolic to continuous metric spaces (e.g., those spaces supporting simple cue weighting and averaging judgment strategies) is often minimal enough that this inaccuracy is tolerable (also see Goldberg, 1968). In short, the benefits of continuity-supported generalization, convergence, and robustness often outweigh the costs of discrete precision.

It is important to recognize the central role of the concept of “robustness” in the present context, due to its intimate relationship to vicarious functioning, Brunswik’s term for the “backbone of stabilized achievement.” Hammond (1996), in commenting on the pioneering work of Dawes and Corrigan (1974) on this issue, has aptly characterized the central insights:

“An interesting and highly important discovery, first introduced to judgment and decision making researchers by R. M. Dawes and B. Corrigan, is that organizing principles of this type [weighting and summing cues] are extremely robust in irreducibly uncertain environments. That is, if (1) the environmental task or situation is not perfectly predictable (uncertain), (2) there are several fallible cues, and (3) the cues are redundant (even slightly), then (4) these organizing principles (it doesn’t matter which one) will provide the subject with a close approximation of the correct inference about the intangible state of the environment, no matter which organizing principle may actually exist therein—that is, even if the organism organizes the information incorrectly relative to the task conditions!” (p. 171)

Thus, intuitive or System 1 cognitive strategies, which are dependent on ecological continuity to support simple weighted averaging, provide reasonable levels of adaptation even if a person has little or no *a priori* knowledge of the structure of the environment (Hammond’s “organizing principle therein”). Also, by supporting generalization and convergence, ecological continuity supports the statistical or neural network learning underlying adaptive behavior.

If we take continuity-supported, intuitive (System 1) cognition as the “backbone of stabilized achievement,” what, then, is the primary functional role of the analytical pole of the intuitive-analytical continuum (Hammond, 1996; Kahneman, 2003)? Current research on dual process theory suggests that discrete, analytical processes seem to be involved in noticing potential anomalies, in monitoring the effectiveness of intuitively-driven behavior, in directing attention to novel events, and in resolving conflicts when multiple and competing intuitive judgments or decisions must be arbitrated (Barnden, 1999; Kahneman, 2003). Activities such as these each touch on the demand for cognition to achieve internal coherence (Hammond, 1996; Thagard, 2000), which is typically understood to be an analytical activity involving the use of discrete logical operations, or System 2 cognition. We should note, however, that Hammond (2000) has recently

proposed the intriguing idea that both coherence and correspondence can each be achieved either intuitively or analytically. For the purposes of this chapter, however, our treatment remains more faithful to previous account offered by Hammond (1996/2000), that:

“The central feature of the correspondence theory of judgment is its emphasis—inherited from Darwin—on the flexibility of the organism in its adaptive efforts, its multiple strategies, its ability to rely on various intersubstitutable features—what are called multiple fallible indicators [i.e., vicarious functioning]” (p. 63).

As such, while possibly simplified, we will assume that the *essence* of correspondence-based achievement lies in the intuitive mechanisms underlying vicarious functioning, whereas the *essence* of coherence-competence lies in analytically dominated, System 2, cognition.

The logical and analytical operations underlying coherence-seeking activities stand in contrast to the demand for cognition to arrive at adaptive solutions corresponding with the facts, constraints, or demands of the external world (empirical accuracy). Understanding this latter, continuity-supported, intuitively-gained, adaptive correspondence between the actual and the perceived was Brunswik’s primary focus. It is our own focus as well, in the sense that we believe that computer and interface technology that short-circuits these adaptive mechanisms will lead to a variety of learning and performance problems in human-automation interaction. For more on this coherence-versus-correspondence (analytical/intuitive) distinction, see Mosier and McCauley (this volume), suggesting that increasingly technological ecologies are indeed placing increased demands on discrete, symbolic, and analytical coherence-seeking (rather than correspondence-seeking) cognition. Note also that Mosier and McCauley’s findings are consistent with what our analysis would expect: People are generally poor at attaining cognitive coherence in discrete, digital ecologies, quite possibly for the reasons we have suggested above.

IMPLICATIONS FOR THE DESIGN OF HUMAN-AUTOMATION INTERACTION

There is ample evidence that people have trouble interacting with discrete interfaces and “navigating” through their many modes, menus, sub, and sub-sub menus (Degani, 2004; Norman, 1999). The frustrations that people have configuring and using the menus of cell phones are popular examples. In more complex systems, it is well documented that digital interactive systems may embody design flaws tied to their discrete mode structures (Degani, 2004). These flaws lead to confusion, and at times, deadly mishaps (Degani and Heymann, 2002).

To better understand these problems, in the following we leverage the distinctions we have made between continuous and discrete ecologies, along with their attendant cognitive implications, with the goal of better supporting human-automation interaction through better interface design. To this point we have emphasized that humans often gravitate toward intuitive (System 1) strategies for making judgments and decisions. Additionally, we have suggested that these strategies are generally robust and reasonably

effective, but are best supported when the human is interacting with a continuous, rather than a solely discrete, ecology (or interface).

We have also pointed out that humans, especially in dynamic and uncertain contexts, often have limited cognitive resources available for effectively navigating the geometry of digital ecologies in an analytical fashion. We, and the findings reported by Mosier and McCauley (this volume) suggest, that everything possible should be done in design to reduce the demands for a system user or operator to rely on coherence-seeking, analytical cognition. If this goal is achieved, then the resources demanded by analytical cognition would be freed-up to for the monitoring, arbitrating, and attending-to-novelty activities to which it appears suited.

Like any cognitive system, a system comprising humans coupled with technology requires both internal coherence and correspondence with the external facts of the world. As described above, we have advocated leaving the attainment of correspondence to the adaptive, intuitive competence of the human operator or user. Regarding the goal of ensuring overall human-automation system coherence, we suggest that this task on should fall largely on the designer, prior to a system being put into operation. With the task of achieving coherence (e.g., ensuring the integrity of the system with respect to efficiency and safety) offloaded to the design process, the human operator or user can then rely to the greatest extent possible on the intuitive mode of cognition known to underlie fluent, robust, and adaptive behavior.

Thus, one can view our solution to the problem of ensuring the goals of both maintaining correspondence with the facts of the world, as well as overall system coherence, in terms of a simple problem decomposition. We advocate assigning the first task to the human operator, or user, who we presume to rely heavily on robust and adaptive intuitive (System 1) cognition, and who is able to benefit by real-time access to timely information. In contrast, our solution assigns the task of ensuring the coherence of overall system operation to the designer, who is shielded from the time pressure and other stressors of real-time operations, and who thus has the cognitive resources available to rely upon analytical tools and analytical (System 2) cognition.

As such, we advocate the use of formal analysis and design techniques to verify that the geometry of human-technology interaction is efficient and safe. This approach to design requires highly detailed, functional analyses of technological artifacts, their interfaces, their environments of use, and the tasks they support. We will first discuss these formal techniques at a general level, and then demonstrate the approach using a variety of concrete examples.

ANALYTICALLY ENSURING COHERENCE IN SYSTEM DESIGN

One formal approach to ensuring the coherence of human-automation interaction in the design phase is a hybrid modeling technique termed Communicating Sequential Processes (“CSP”—see Hoare, 1985; Schneider, 2000). Techniques such as CSP ensure the overall integrity, or coherence, of systems comprised of multiple computational elements to, for example, guard against deadlocking, the system entering dangerous or “illegal” states, and so on. Relatedly, a variety of discrete, finite-state-machine models

can also be used to create formal representations of human interaction with control systems. A basic element of such a description is a labeled, directed graph which captures system states, events, conditions, and transitions. For example, while an aircraft is in “Cruise” mode (current state), and button x is pressed, and the descent profile is armed, (two conditions that when TRUE trigger an event), the aircraft control system will transition to “Descent” mode (a new state).

Due to the lack of continuity in discrete space, human operators or users cannot always reliably anticipate the future mode configuration of a machine unless they have a detailed internal model of the machine’s behavior, including its states, conditions, and transition logic. (Degani, Shafto, & Kirlik, 1999). They must rely on *a priori* knowledge (e.g., an internal model), because lacking a continuous, metric representation of the technology, they cannot deploy their robust, intuitive cognition in the sense discussed by Dawes, Hammond, and Kahneman.

Human interaction with automation is currently evaluated through extensive simulation. Smith et al. (1998) note the serious problems that can occur with solely simulation-based, (empirical) evaluation of complex human-automation interaction. Because of their combinatorial nature, discrete and hybrid systems have enormous state spaces, and empirical techniques that necessarily sample only a small subset of the entire state space are clearly insufficient as a basis for design verification. As such, predictive, model-based, formal methodologies for design verification are critical for identifying deficiencies early in the design phase. To verify system safety, a safety specification is initially represented as a restricted region of the state space in which the system should remain. On the basis of this functional analysis, the goal is to synthesize automation design that guarantees that the state of the system will remain within a safe region.

Using the same modeling formalisms used to describe machine behavior (the “machine model”), it is possible to also describe the information provided to the user. This “user model” of the system, which is based on the information provided to the user (e.g., pilot) via displays, manuals, procedures, training, and personal experience, may differ from the actual “machine model” of the system. With these two models in place, it becomes possible to systematically and comprehensively verify that the interfaces (and all other information provided to the user) are correct. This, in principle, is done by constructing a *composite* model where the user-model states and machine-model states are combined into state-pairs. Next, the verification process simulates an activation of the composite model, where the user-model and the machine model evolve concurrently in a synchronized manner (Heymann and Degani, 2002).

This verification process detects three types of user-interface inadequacies that are based on the criteria set by Heymann and Degani. The first inadequacy—the existence of *error* states—occurs when the user interface indicates that the machine is in one mode, when in fact the machine is in another. Interfaces with error states lead to faulty interaction and errors. The second inadequacy—the existence of *restricting* states—occurs when the user is unaware that certain user interactions can trigger additional mode changes in the machine. Interfaces with restricting states tend to surprise and confuse users. The third inadequacy—the existence of *augmenting* states—occurs when the user is informed that a certain mode is available, when in fact the machine does not have that mode, or access to

certain modes is disabled. Interfaces with augmenting states puzzle users. Other problems can also be revealed by such analytical, formal analysis. For example, if the overall composition model “deadlocks,” the system may have reached a state in which no effective control input is available to ensure system safety.

EXPLORING THE GEOMETRY OF HUMAN-TECHNOLOGY INTERACTION: CONCRETE EXAMPLES

To illustrate our perspective on designing to enhance human-technology interaction, we now turn our attention to three everyday examples:

Pedestrian Crossing Signals

The first example consists of a pedestrian crossing signal found at many intersections with traffic lights. The most familiar ones have three states (or modes): a red hand for “DON’T WALK,” a walking-person symbol for “WALK,” and a flashing hand to indicate that the light is about to turn red and crossing is unsafe. Newer models, which now can be found in many intersections, provide, in addition to the flashing hand, a digital countdown display (in seconds), indicating when crossing will become unsafe (and the “DON’T WALK” symbol will appear).

The old interface is discrete: don’t walk, walk, hurry up. The new interface also indicates the time remaining until crossing will become unsafe (the “DON’T WALK” symbol appears). The new design is a “hybrid interface,” containing both discrete states and one close approximation to a continuous variable (the “countdown time”). Most people prefer the new interface because it reduces their fear that the light will turn red while they are in the middle of an intersection. When the light begins to flash and the countdown is displayed, the walker can assess the situation, and decide if he or she needs to run, walk, or perhaps turn back and wait for the next light. Note that by augmenting the interface with continuous information, better supporting the walker’s intuitive judgment, the entire human-technology system gains a greater level of coherence (in this case, through an increased level of assurance that safety goals are met).

Traffic Signals

Consider the next scenario: now you are in your car driving toward an intersection, and you can see a green traffic light in the distance. As you approach the intersection, the light turns yellow. What should you do? Brake or proceed through? Figure 2 shows a region in which braking is safe; i.e., given any combination of speed and distance within the dark gray region, the car will stop before the intersection. This region is determined by calculating the stopping distance for every speed from 1-60 miles per hour based on a car’s maximum braking performance. The yellow light is 4 seconds long, and we take into account that it takes the average driver 1.5 seconds to react and press on the brakes.




Figure 2. The Safe Braking Region (X-Axis is Distance from Intersection, Y-Axis is Speed)

As for the other option, driving through the intersection, let us assume that the driver, once he or she observes the yellow light, simply maintains the current speed. Figure 3 depicts the safe drive-through region (while maintaining constant speed). Now that we know the consequences of either stopping before the intersection or driving through it, we can turn to the decision itself. So how do we know, when the yellow light appears, what to do? The truth is that in most cases we don't. We don't have the analytically-derived graphs of Figures 2 and 3 available, and therefore our decision is based on our intuitive estimation of what Gibson and Crooks (1938) called "the field of safe travel." Most of the time our intuition works, but in some cases we either pass through a red light or try to brake but end up stopped in the intersection.




Figure 3. The Safe Proceed-Through Region

Now, consider designing a dashboard interface to help the driver make this decision. To do so, we first need to map the physical space and the system behavior *vis-à-vis* the operational demands (stopping or proceeding through). We do this by considering the “safe braking” region and the “safe proceed-through” region together. Figure 4 shows that the composition of the two graphs divides the physical space into four sub-regions. The dark gray region is “safe braking,” the light gray is “safe proceed-through,” and the black region is where these overlap (if you are in the black region, you can either brake or proceed through, and you'll still be safe and legal).




Figure 4. Composite of the Safe-Breaking and Safe Proceed-Through Regions

Unlike the above three regions from which a safe control action can be taken, the fourth (hatched) region represents combinations of speed and distance from which you *cannot* safely brake or proceed safely: If you try to stop, even at maximum braking force, you will find yourself entering the intersection on red. And if you proceed through, you will reach the white line with a red light above you. In other words, if you are in the hatched region when the light turns yellow, you *will* commit a violation, no matter what you do (and as for the tempting option of “gunning it” through, calculations show that acceleration only reduces the size of the hatched region, but does not eliminate it). See Oishi, Tomlin, & Degani, 2003 for the technical details of this kind of analysis.

This problematic region, which is well known to traffic engineers, is called the “dilemma zone” and exists in many intersections (Liu, Herman, & Gazis, 1996). From a design perspective, this implies that we not only need to display if the one in the “brake” or “drive through” region, but also the “dilemma zone.” But if we think about it for a minute, just providing a discrete indication, “You are in a dilemma zone,” is not enough. Why? Because it’s too late! For the interface to ensure the coherent, safe operation of the entire human-technology system, it must provide a *continuous* indication of the proximity of the car to the dilemma zone, so that the driver can avoid entering it. And unlike the pedestrian light, where the countdown timer is a welcome, “nice to have,” addition, in the yellow light case, it is an imperative. An interface that provides solely discrete indications (“brake,” “drive through,” “in the dilemma zone”) is unsafe!

Automatic Landing Systems

Our final example concerns one aspect of the automatic flight control system on board modern airliners designed to make a fully automatic landing. An automatic landing system, or “autoland,” is commonly used in bad weather, specifically in a condition called “zero-zero” in pilot lingo, meaning visibility is zero and the clouds or fog reach all the way to the ground.

In these severe conditions, only the autoland system is permitted to make a landing. But there is one option that is always available to the pilot—to discontinue the approach and abort the landing, and then command the aircraft to climb out. Such aborted landings, or “go-arounds,” are well-practiced maneuvers, with the goal of taking the aircraft away from an unsafe situation such as an autoland component failure, electrical system failure, or any other malfunction making the approach and landing unsafe. In some cases, a go-around is requested by air traffic control. For example, an aircraft may come too close to another aircraft on approach, or there may be debris, a vehicle, or another aircraft on the runway on which the aircraft is to land.

We wish to consider the kind of information that must be provided to the pilot while the autoland system is making the approach and landing. In particular, our focus is on the critical, last 60 feet of the approach. Just as we did in the case of the yellow light, we first need to map out the physical space and the technological (control system) spaces. Then we need to relate them to the two control options available: either the airplane lands or attempts a go-around.




Figure 5. Safe Landing

The funnel shape in Figure 5 depicts the region from which an autoland system can make a safe landing. The three axes are the aircraft's altitude above the runway, speed, and flight-path angle (which is the angle at which the airplane descends towards the ground—see top of Figure 5). In principle, the shape is computed in the same way as the “safe breaking” and “safe proceed through” regions in the yellow light example. We start from touchdown, where the flight-path angle should be between 0 and -2 degrees, and work our way back. (If the angle is greater than zero, the airplane will not be able to land; if the angle is less than -2 degrees, the aircraft's tail will hit the ground). For each altitude, from 0 to 60 feet, we compute the speed and flight-path angle the autopilot needs to maintain, such that eventually the aircraft *will* make a safe landing.




Figure 6. Safe Go-Around.

As for the second option, the “safe go-around,” Figure 6 depicts a rather large region, because unlike the safe landing region that funnels down to the runway at a tightly constrained angle and speed, the go-around can be executed safely at a variety of flight-path angles and speeds (see Oishi, Tomlin, & Degani, 2003 for the technical details of these computations).

We now combine the safe landing region and the safe go-around region (just as we did in the yellow light example). For the sake of illustration, consider what will happen when a go-around is commanded when the aircraft is 20 feet above the ground. Figure 7 is a slice of the safe-landing region at this altitude. Figure 8 is a similar graph for the safe go-around region.




Figure 7. Safe Landing Region at 20 Feet




Figure 8. Safe Go-Around Region at 20 Feet

The composite graph of Figure 9 shows three emerging sub-regions: the gray region is the safe go-around region, the black region depicts where safe go-around overlaps with safe-landing, and the hatched region is solely the safe-landing region. Because a go-around is a situation that can occur at any time during landing (and there may be no advance warning of when the maneuver will be needed), the black region is where we always want to be: from here the autopilot can make a safe landing, and if a go-around is needed, it can be executed safely by the pilot.




Figure 9. Composite Graph at 20 Feet.

The hatched area, however, is where we *don't* want to be. Under normal conditions, the autopilot will try to make the landing when the flight-path angle is close to 0 degrees, but under less than nominal conditions, such as gusts or a strong tailwind, the autopilot may be operating in the hatched region in which, nevertheless, a safe landing can be completed. Safe landing, however, is only *one* requirement. We also need to be in the “safe go-around,” region, and herein lies the problem: If the autopilot is operating in the hatched region, the pilot will not be able to execute a safe go-around, either at this altitude of 20 feet, or at any lower altitude. If the autopilot is operating in the hatched region and a go-around is attempted, the aircraft may stall!

Lessons Learned

To summarize, the yellow light and autoland examples highlight several important issues in the design of automated control systems in general, and human-automation interaction in particular. One issue has to do with the complexity of the physical space and system behavior that produces the “dilemma zone” and the “Safe Landing/Unsafe Go Around” regions. Identifying these regions requires analysis, which can yield counter-intuitive results (did you know that when you get a citation for a red light violation, in some cases it may be not your fault?). This analysis requires a perspective drawing on the discrete-versus-continuous ecology distinction, and analytical tools for mapping out the geometry of human-technology interaction.

Another issue, centering on interface design, concerns a tendency many designers have to provide solely discrete indications. There are two problems with this. First, a driver using such a discrete alerting interface has no warning as to when the “dilemma zone” is being neared. The interface would just “light up” or alarm. And if it does, it’s too late, because the human-technology system is already in an unsafe region. Likewise, by providing the pilot with solely a discrete indication of the “unsafe go-around” region, we haven’t solved the problem. Since, in most automated control systems, the existence of unsafe

regions is inevitable, interfaces must also provide the user with continual guidance as to system's proximity to the unsafe region.

Finally, an important distinction should be made between the driving example and the aviation example. In the yellow light case, it is the driver who (manually) enters the unsafe region from which any action (brake or proceed-through) will have dangerous consequences. In contrast, in the autoland case, it is the *automation* that takes the user into an unsafe region, from which recovery is difficult, and at times, impossible. This situation, arising from the complex, and not yet well understood, interaction between discrete modes and continuous processes is termed "automation lock." Such locks have lulled many pilots into dangerous situations, some unfortunately resulting in accidents (FAA, 2004; Degani, 2004: Chapter 17).

CONCLUSIONS

The title of this chapter is borrowed from Bryant (1985, pp. 87), who asks the question, "What makes analysis work?" His answer is built around the concept of *continuity*. Continuity makes it possible to represent and reason about complex, dynamic, and non-repeating phenomena. It makes available a range of tractable methods for interpolation, extrapolation, approximation, and the integration of fragmentary data. All are essential to vicarious functioning.

The products of the psychological processes of perception, cognition, and learning can be described in terms of cognitive representations, ranging from the continuous to the symbolic and discrete. Basic perceptual processes enforce assumptions about spatio-temporal continuity (Barlow, 2001; Shepard, 1999). And even in situations where discrete, symbolic representations are more precise, continuity is generated to support adaptive, heuristic reasoning and learning in abstract symbolic spaces (Hastie & Dawes, 2001; Kahneman, 2003; Landauer et al., 1997).

Digital control systems, and their corresponding discrete (mode) interfaces, disrupt continuity. In the extreme case, we are left with a step-by-step search through a space in which each move is equally capable of taking us into a familiar room or a blind alley. Furthermore, it is very difficult (and most of the time impossible) to design or learn to operate large and complex systems that are based on unrestricted, discrete spaces. Designers are overwhelmed by the large state spaces of technological systems (which may include thousands, if not millions, of possible states). And users are faced with the challenge that achieving an exact understanding of a system's behavior is impossible, yet reliable abstractions or approximations are unavailable.

Egon Brunswik emphasized the role of pattern-pattern correlations in complex perception and action. This is the fundamental principle underlying the lens model. In natural environments, spatial cues and spatio-temporal trajectories rarely or never repeat themselves exactly. An underlying continuity must be assumed to enable inference and generalization from fragmentary, non-repeating, data to stable objects and predictable trajectories. Vicarious functioning, or purposive behavior, is essentially adaptive. Local adaptations involve corrections and alternate paths. Global adaptations involve learning

converging to stable pattern-pattern relations, in that the patterns are more reliable and robust than single cues or individual, cue-object associations.

Early in this chapter, we described how canonical correlation analysis can be used to identify pattern-pattern correlations in the way pilots interact with automation, and, more importantly, how to capitalize on the regularities of these patterns to identify unusual deviations and dangerous outliers. It may be appealing to think, that because stochastic patterns repeat with some regularity even in digital systems, that human users or operators can adapt to them just as they can adapt to the familiar, physical environment. This, however, is only true as long as the system remains in a restricted, familiar subspace of the overall state space. Problems arise from the fact that, although approximations to continuous processes can have bounded errors, continuous approximations to discrete systems are not so well-behaved. Outside the nominal, familiar subspace in which a system nominally operates, that is, in unusual situations and at unexpected times, underlying ecological discontinuities may suddenly manifest themselves, and the resulting consequences of the errors in a continuous approximation may be arbitrarily large.

Currently, a tendency exists among designers of both everyday devices and more complex systems to mimic the underlying, discrete nature of computer-based artifacts with a “simplified,” all-discrete interface. Since many consumer devices and all automated control systems are hybrid (i.e., they contain both discrete and continuous process), solely discrete interfaces to automation (e.g., mode settings, alerts, alarms) may abstract away important, and sometimes, even critical information.

Due to the paucity of formal, rigorous, and systematic methods for supporting human-technology interaction design, the tasks of creating these abstractions and their attendant interface design simplifications are too often performed intuitively, and in an ad-hoc way. The result, in many cases, is user frustration, confusion and, in the case of high-risk systems, the possibility of disaster. We hope that the insights we have offered here, advocating the formal analysis of the geometry of human-technology interaction, will take at least a small step toward remedying this situation.

ACKNOWLEDGMENTS

The work described in this chapter was conducted as part of NASA’s Intelligent Systems program with the intent of supporting the development of future interfaces for controlling and managing semi-autonomous rovers. The research underlying the cockpit automation examples provided in this chapter was supported by NASA’s Aerospace Technology research programs. The first author would like to thank Lisa Burnett for her teaching of flow and continuity and Yosi Amaram for a fruitful discussion that contributed to the writing of this chapter.

REFERENCES

- Barlow, H. (2001). *The exploitation of regularities in the environment by the brain*.
<http://www.bbsonline.org/Preprints/OldArchive/bbs.barlow.html>
- Barnden, J.A. (2001). Uncertainty and Conflict Handling in the ATT-Meta Context-based System for Metaphorical Reasoning. In *Proceedings of CONTEXT'01: Third International Conference on Modeling and Using Context*, Dundee, Scotland.
- Bruner, J. S. (1973). *Beyond the Information Given: Studies in the Psychology of Knowing*. New York: Norton.
- Brunswik, E. (1950, 1952). *The Conceptual Framework of Psychology*. Chicago: University of Chicago Press.
- Brunswik, E. (1956). *Perception and the Representative Design of Psychological Experiments*. Berkeley, CA: University of California Press.
- Bryant, V. (1985). *Metric Spaces: Iteration and Application*. Cambridge, UK: Cambridge University Press.
- Bullock, S. & Todd, P. M. (1999). Made to measure: Ecological rationality in structured environments. *Minds and Machines*, 9, 497-541.
- Casner, S. M. (1994). Understanding the determinants of problem-solving behavior in a complex environment. *Human Factors*, 36(4), 580-598.
- Cliff, N. (1987). *Analyzing Multivariate Data*. San Diego: Harcourt Brace Jovanovich.
- Connolly, T. (1999). Action as a fast and frugal heuristic. *Minds and Machines*, 9, 479-496
- Court of Inquiry. (1992). Report on accident to Indian Airlines Airbus A-320 aircraft VT-EPN at Bangalore on 14th February 1990. Indian Government.
- Dawes, R. M. (1979). The robust beauty of improper linear models. *American Psychologist*, 34, 571-582.
- Degani, A. (2004). *Taming HAL: Designing interfaces beyond 2001*. New York: Palgrave Macmillan.
- Degani, A. (1996). *Modeling human-machine systems: On modes, error, and patterns of interaction*. Unpublished doctoral dissertation. Atlanta, GA: Georgia Institute of Technology.
- Degani A, & Heymann, M. (2002). Formal verification of human-automation interaction. *Human Factors*, 44(1), 28-43.
- Degani, A., Shafto, M., and Kirlik, A. (1999). Modes in human-machine systems: Constructs, representation, and classification. *International Journal of Aviation Psychology*, 9(2), 125-138.
- De Leeuw, J. & Michailidis, G. (2000). Graph layout techniques and multidimensional data analysis. In F.T. Bruss and L. Le Cam (Eds.), *Game Theory, Optimal Stopping, Probability and Statistics. Papers in honor of Thomas S. Ferguson*, IMS Lecture Notes-Monograph Series, 219-248.
- Federal Aviation Administration (FAA) (2004). Advisory Circular to Supplement Federal Aviation Regulation 25.1329: Automatic Pilot Systems. Washington, DC.

- Freed, M., & Remington, R.W. (1998). A conceptual framework for predicting errors in complex human-machine environments. In *Proceedings of the 1998 Conference of the Cognitive Science Society*, University of Wisconsin, Madison, WI.
- Gibson, J.J. & Crooks, L.E. (1938). A theoretical field-analysis of automobile driving. *American Journal of Psychology*, 51, 453-471.
- Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650-670.
- Goldberg, L. R. (1968). Simple models or simple processes? Some research on clinical judgment. *American Psychologist*, 23, 483-496.
- Hammond, K. R. (1996). *Human Judgment and Social Policy*. New York: Oxford University Press.
- Hammond, K. R. (2000). *Judgments Under Stress*. New York: Oxford University Press.
- Hammond, K. R. (2000/1996). Coherence and correspondence theories in judgment and decision making. In T. Connolly, H. R. Arkes, and K. R. Hammond (Eds.), *Judgment and Decision Making: An Interdisciplinary Reader*, 2nd Edition. (pp. 54-65). Cambridge, UK: Cambridge University Press. Chapter originally appeared in Hammond (1996, op cit.)
- Hastie, R., & Dawes, R.M. (2001). *Rational choice in an uncertain world: The psychology of judgment and decision making*. Thousand Oaks, CA: Sage.
- Heymann, M. and Degani, A., (2002). *On abstractions and simplifications in the design of human-automation interfaces*. NASA Technical Memorandum 2002-21397, Moffett Field, CA.
- Hoare, C. A. R. (1985): *Communicating Sequential Processes*, New York: Prentice Hall.
- Hubbard, T.L. (1999). How Consequences of Physical Principles Influence Mental Representation: The Environmental Invariants Hypothesis. In P.R. Killeen & W.R. Uttal (Eds.), *Fechner Day 99: The End of 20th Century Psychophysics. Proceedings of the Fifteenth Annual Meeting of the International Society for Psychophysics*. Tempe, AZ: The International Society for Psychophysics.
- Jha, P., & Bisantz, A.M. (2001). Modeling Fault Diagnosis in a Dynamic Process Control Task Using a Multivariate Lens Model. In *Proceedings of the HFES 45th Annual Meeting*, Minneapolis/St. Paul, MN.
- Kahneman, D. (2003). A Perspective on Judgment and Choice: Mapping Bounded Rationality. *American Psychologist*, 58(9), 697-720.
- Kirlik, A. (in press). Work in progress: Reinventing intelligence for an invented world. In R. J. Sternberg & D. Preiss (Eds.), *Intelligence and Technology*. Mahwah: NJ: Erlbaum.
- Kirlik, A. (2001). On Gibson's Review of Brunswik. In K.R. Hammond & T.R. Stewart (Eds.), *The Essential Brunswik: Beginnings, Explications, Applications*. New York: Oxford University Press.
- Kirlik, A. (1995). Requirements for psychological models to support design: Towards ecological task analysis. In, J. Flach, P. Hancock, J. Caird & K. Vicente (Eds.). *Global Perspectives on the Ecology of Human-Machine Systems* (pp. 68-120). Mahwah, NJ: Erlbaum.

- Klein, G. A. (1999). *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press.
- Landauer, T. K., Laham, D., Rehder, B., & Schreiner, M. E., (1997). How well can passage meaning be derived without using word order? A comparison of latent semantic analysis and humans. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the 19th Annual Meeting of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.
- Liu, C., Herman, R. and Gazis, D. (1996). A review of the yellow interval dilemma. *Transportation Research A*, 30(5), 333–348.
- Lowe, G. (1992). *Some extensions to the probabilistic, biased model of Timed CSP*. Technical Report PRG-TR-9-92. Oxford University Computing Laboratory.
- Lund, K. & Burgess, C. (1996). Producing High Dimensional Semantic Spaces from Lexical Co-occurrence. *Behavior Research Methods, Instruments and Computers*, 28, 203-208.
- McGreevy, M.W. & Statler, I.C. (1998). *Rating the Relevance of QUORUM-Selected ASRS Incident Narratives to a "Controlled Flight into Terrain" Accident*. NASA Technical Memorandum 1998-208749. Moffett Field, CA: NASA-Ames Research Center.
- Ministry of Planning, Housing, Transport and Maritime Affairs (1989). Investigation commission concerning the accident which occurred on June 26th 1988 at Mulhouse-Habsheim to Airbus A-320, registration number F-GKC.
- Norenzayan, A., Smith, E.E., Kim, B.J., & Nisbett, R.E. (2002). Cultural Preferences for Formal versus Intuitive Reasoning. *Cognitive Science*, 26, 653-684.
- Norman D. (1999). *The Invisible Computer*. Boston: MIT Press
- Reason, J. (1990). *Human Error*. Cambridge, UK: Cambridge University Press.
- Rouse, W. B. (1983). Models of human problem solving: Detection, diagnosis, and compensation for system failures. *Automatica*, Vol. 9, pp 613-625.
- Roscoe, A.W. (1998). *The Theory and Practice of Concurrency*. London: Prentice Hall.
- Shafto, M., Degani, A., and Kirlik, A. (1997). Canonical Correlation Analysis of Data on Human-Automation Interaction. *Proceedings of the 41st Annual Meeting of the Human Factors and Ergonomics Society*. Albuquerque, NM.
- Shepard, R.N. (1999). *Perception, imagery, and science*. Hitchcock Lecture, University of California, Santa Barbara.
- Sheridan, T. (1997). Supervisory control. In G. Salvendy (Ed.), *Handbook of Human Factors*. New York: Wiley.
- Smith, S., Duke, D.J., Marsh, T., Harrison, M.D., & Wright, P.C. (1998). *Modeling Interaction in Virtual Environments*, UK-VRSIG '98, Exeter, UK.
- Sun, R. (2002). *The Duality of Mind: A Bottom-Up Approach Toward Cognition*. Mahwah, NJ: Lawrence Erlbaum.
- Thagard, P. (2000). *Coherence in Thought and Action*. Cambridge, MA: MIT Press.
- Wolf, B. (May, 1999). Vicarious functioning as a central process-characteristic of human behavior. In K.R. Fischer & F. Stadler (Eds.), "*Wahrnehmung und Gegenstandswelt*". *Zum Lebenswerk von Egon Brunswik*. Wien: Springer. University of Landau, Center of Educational Research. <http://www.brunswik.org/notes/essay4.html>

Zeigler, B.P., Praehofer, H., & Kim T.G. (2000). *Theory of Modeling and Simulation: Integrating Discrete Event and Continuous Complex Dynamic Systems* (second edition). San Diego: Academic Press.